INTERNAL AUDIT REPORT

Vector Consciousness Prediction (VCP) — Full Project Review

The Self-Experiencing Universe: A Manifesto for an Emergent Reality (Lis, 2025)

Waldemar Lis (ontology) | Claude / Anthropic (formal analysis) — March 2026

Full archive: https://archive.org/details/@not_foryou

> **AUDIT STATUS:** *This audit was conducted in March 2026, prior to any responses from researchers currently in correspondence (Prof. Lars Chittka, Queen Mary University of London; Prof. Nicola Clayton, University of Cambridge). It constitutes a pre-external-feedback internal review. All findings documented here precede any external scientific input.*

---

1. Scope and Method

This document reports the results of a systematic internal review of the full Vector Consciousness Prediction project, covering the following materials:

- The Self-Experiencing Universe: A Manifesto for an Emergent Reality (Lis, 2025) — foundational ontology
- Eight structural isomorphism simulations (Coherence Study, January–February 2026)
- HoTT Formalization v1 and Revision v2 (February 2026)
- Research Protocol — First Falsifiable Prediction of Relational Ontology (February 2026)
- Literature Review — Empirical Verification of VCP (February 2026)
- VCP Simulations 1–4 (February–March 2026)
- Sepia officinalis Opt-Out Prediction — Full Technical Report and Public Summary (March 2026)

The audit applied the following standard: identify logical errors, methodological asymmetries, unsupported claims, and any points at which a scientific reviewer could mount a principled objection. Confirmation that something works was treated as equally important as identification of a flaw. No results were adjusted retroactively.

---

2. What Holds

The following structural elements were reviewed and found to be internally consistent and defensible:

2.1 The Logical Chain

The progression from foundational ontology through mathematical formalization to empirical prediction is coherent throughout. Each transition is explicitly motivated. No step assumes a conclusion that the preceding step did not establish.

2.2 Withdrawal of AQE and Revision to IT-Region

The decision to withdraw the Axiom of Qualitative Emergence (AQE) in HoTT v2 and replace the binary Integration Threshold with a topological region in (s × i × r) space is correctly motivated by naturalistic observation and correctly executed formally. The revised axiom set HoTT-RO v2 = HoTT + {UA, AFM} is parsimonious and coherent. The Gödelian limit of Relational Potential (RP) is correctly identified as an inherent structural feature, not a gap.

2.3 Numerical Calculations

All calculations in Simulations 1–4 were independently verified. The key geometric finding — d(bee_ind, Sepia) = 0.3265 < d(bee_ind, swarm) = 0.6226, ratio 1.907 — survives full empirical grounding. The recovery check r_recovered = √(0.9800 × 0.5740) = 0.7500 is exact. No arithmetic errors were found.

2.4 Falsification Conditions

Falsification conditions are explicitly and precisely stated in the Research Protocol and the Sepia technical report. The prediction is falsified if opt-out on difficult trials falls below 0.60 (bee level). The prediction is partially confirmed but requires upward revision of r_sepia if the result exceeds 0.92. Both conditions are operationally testable.

2.5 Exclusion of Propensity

The exclusion of Propensity from the active VCP model is consistent throughout all documents. No document in the series reintroduces it or relies on it.

2.6 Literature Review Findings

The behavioral fingerprints identified in the literature review (Schnell et al. 2021, Billard et al. 2020, Seeley 2003/2004/2010) are correctly interpreted through the VCP framework. The contrast between individual-level rapid adaptation in Sepia and network-level quorum-dependent adaptation in the swarm is structurally precise and consistent with the predicted vector profiles.

---

3. Flags — Items Requiring Attention

Four items were identified that require attention before public dissemination. They are ordered by importance.

⚠ FLAG 1 — SERIOUS r_sepia — Sensitivity Analysis Absent

r_bee = 0.4714 is derived from a formal quantitative experiment (Perry & Barron 2013: measurable opt-out rates). r_sepia = 0.75 is derived from qualitative behavioral observations compiled by Prof. Jennifer Mather (three observations: unseen arm monitoring, prospective arm substitution, head bob as epistemic check). These are two different epistemological levels. The full Sepia opt-out prediction of 0.836 is directly anchored to r_sepia = 0.75. The technical report conducts sensitivity analysis for opt-out_easy and opt-out_impossible — but not for r_sepia itself, which is the primary anchor. A scientific reviewer will identify this asymmetry immediately.

RECOMMENDED ACTION: Add a dedicated section to the Sepia technical report showing how the opt-out difficult prediction shifts across r_sepia = {0.65, 0.70, 0.75, 0.80, 0.85}. This does not weaken the theory — it demonstrates controlled awareness of the uncertainty and strengthens the document's scientific credibility. The asymmetry exists; the question is whether it is documented or left for a reviewer to find.

◆ FLAG 2 — MODERATE Young (1971) — Species Substitution Unjustified in Text

Simulation 4 uses Young, J.Z. (1971), The Anatomy of the Nervous System of Octopus vulgaris, as the first empirical source for i_sepia (integration in Sepia officinalis). This is anatomical data from Octopus vulgaris, not Sepia. The substitution is phylogenetically reasonable — both are coleoid cephalopods with comparable central-ganglion-to-peripheral-neuron architecture — but this reasoning is not stated anywhere in the document. A reviewer familiar with cephalopod neurobiology will notice the species mismatch and may question the derivation.

RECOMMENDED ACTION: Add a single sentence of explicit justification: 'Young (1971) documents the coleoid cephalopod architecture in which approximately two-thirds of neurons are distributed in peripheral ganglia rather than the central brain. This organizational principle is conserved across Octopus and Sepia (Hanlon & Messenger 1996), justifying its use as a source for i_sepia.' Alternatively, replace with a Sepia-specific anatomical source if one is identified.

◇ FLAG 3 — MINOR Research Protocol — phi Comparability Assumption Is Outdated

The Research Protocol (February 2026) identifies the comparability of phi(swarm) ≈ phi(cuttlefish) as 'an open technical gap requiring empirical verification.' This gap was resolved in subsequent analysis: the IIT 4.0 Exclusion Principle establishes phi ≈ 0 for spatially distributed collective systems on principled (not merely technical) grounds. The slow communication timescale of swarm coordination (seconds to minutes vs. milliseconds in neural systems) means the IIT exclusion is structural. VCP now operates as an independent framework that does not depend on phi comparability. A reader encountering the Research Protocol without knowledge of the subsequent analysis will believe VCP rests on an unverified assumption.

RECOMMENDED ACTION: Add a brief dated addendum at the end of the Research Protocol: 'Addendum, March 2026: Subsequent analysis of IIT 4.0 (Tononi et al.) confirmed that phi ≈ 0 for swarm-type distributed collectives on principled grounds (Exclusion Principle combined with communication timescale constraint). VCP accordingly operates independently of phi comparability. Condition alpha of the protocol is superseded by this finding.'

○ FLAG 4 — COSMETIC Coherence Study — Unsupported Probability Claim

The Coherence Study contains the statement: 'The probability of obtaining this gradient by chance is very low.' No statistical calculation supports this claim. The data that precede this sentence are already compelling on their own: seven domains, a structured coherence gradient that is directionally predicted by the ontology's commitments. The unsupported probability claim is weaker than the evidence it is trying to characterize, and it is the kind of statement that invites methodological criticism.

RECOMMENDED ACTION: Remove the sentence entirely, or replace with: 'The gradient is structurally predictable from the ontology — domains sharing its structural commitments (relational co-constitution without substance presuppositions) score higher; domains assuming time-as-primitive or substance-adjacent structures score lower. This predictability is itself a result.'

4. Summary Table

Element	Status	Priority
Internal logic chain (book → simulations → HoTT → VCP)	HOLDS	—
Numerical calculations (Simulations 1–4)	CORRECT	—
Falsification conditions	CLEARLY STATED	—
r_sepia — sensitivity analysis	MISSING	HIGH
Young (1971) — species substitution justification	UNJUSTIFIED (implicit)	MEDIUM
Research Protocol — phi comparability assumption	OUTDATED	LOW
Coherence gradient probability claim	UNSUPPORTED CLAIM	LOW

---

5. Overall Verdict

The VCP project is structurally sound. The logical chain from foundational ontology through mathematical formalization to empirical prediction is internally consistent, numerically correct, and produces a falsifiable prediction with clearly stated falsification conditions.

None of the four flags identified in this audit constitutes a structural error or an invalidating gap. Flags 1 and 2 are asymmetries that a scientific reviewer will notice — and that are better addressed proactively than reactively. Flags 3 and 4 are corrections that require minimal effort and improve document clarity.

The central claim of the project — that the individual honeybee is geometrically closer in VCP space to Sepia officinalis than to its own swarm, and that this geometry generates a falsifiable prediction of opt-out behavior in an experiment not yet conducted — survives full internal audit.

> *The framework is waiting for experimental confirmation. The internal work is complete.*

Waldemar Lis — The Self-Experiencing Universe Project

Internal audit conducted: March 2026

Prior to external researcher responses

https://archive.org/details/@not_foryou